arXiv:1401.7886v2 [cs.DS] 13 Jun 2014

Balancing lists: a proof pearl*

Guyslain Naves[1] and Arnaud Spiwack[2]

[1] Aix Marseille Université, CNRS, LIF UMR 7279, 13288, Marseille, France
guyslain.naves@lif.univ-mrs.fr
[2] Inria Paris-Rocquencourt
Ens, Paris, France
arnaud@spiwack.net

Abstract. Starting with an algorithm to turn lists into full trees which uses non-obvious invariants and partial functions, we progressively encode the invariants in the types of the data, removing most of the burden of a correctness proof.
The invariants are encoded using non-uniform inductive types which parallel numerical representations in a style advertised by Okasaki, and a small amount of dependent types.

1 Introduction

Starting with a list lst, we want to turn it into a binary tree tr of the following form (in Ocaml):

```
type α tree =
  | Node of α tree * α * α tree
  | Leaf
```

With the constraints that lst must be the infix traversal of tr and that tr must be *full*, in the sense that every level except the last are required to be completely filled. Such a function turns, in particular, sorted lists into balanced binary search trees.

There are a number of folklore algorithms to achieve this result in linear time. Here we consider one of these algorithms, presented in Section 2, which repeatedly pairs up trees of height h in a list to form a list of trees of height $h+1$. Our interest in this algorithm sprouts from the fact that its correctness is not obvious; in particular the invariants are complex: the main loop operates on a list of length $2^k - 1$ whose elements are alternately of two distinct forms.

In Sections 3 and 4 we show refinements of the algorithm where the invariants are pushed into the types, leading to a complete and short proof of correctness in Coq.

* This research has received funding from the European Research Council under the FP7 grant agreement 278673, Project MemCAD

2 A balancing algorithm

We start by giving a first, simple, implementation of the balancing algorithm. The heart of the algorithm relies on using an alternating list of length $2^k - 1$, where odd-position elements are trees and even-position elements are labels, of type α (indices starting from 1). A full tree of height k can be decomposed into the first $k-1$ levels, containing $2^{k-1}-1$ internal nodes, and the kth level, which contains both nodes and leaves. Thus, the $2^{k-1}-1$ labels in the alternating list will be used to label the internal nodes in the $k-1$ first levels of the balanced tree, while the 2^{k-1} trees, all of height at most one at first, will constitute the level k.

Though we could encode labels as trees of height one in the alternating list, we rather use an appropriate type for the sake of readability:

```
type α tree_or_elt =
  Elt of α
| Tree of α tree
```

We decompose the problem into two parts: computing an alternating list of length $2^k - 1$ from an arbitrary list of labels, and then transforming this alternating list into a balanced tree. We first show how to solve the second part: turning an alternating list into a full tree.

Given an alternating list lst, by pairing the trees in lst using only one traversal of the list, we obtain an alternating list with exactly half as many trees. Each pairing requires two trees and one label used as a root. In order to build a list that is alternated, we also need a second label, that is kept as a single element. This explains why we consider at each step the four first elements of the list.

A single traversal, encoded by pass : α tree_or_elt list → α tree_or_elt_list, reduces an alternating list of length $2^k - 1 \geqslant 3$ to an alternating list of length $2^{k-1} - 1$. By iterating this process using loop : α tree_or_elt list → α tree, we reduce the original list to a list of length one, whose one element is a balanced tree t such that the infix traversal of t is the initial list.

```
let join left node right = Tree (Node (left, node, right))

let rec pass = function
  | Tree left :: Elt root :: Tree right :: Elt e :: others →
      join left root right :: Elt e :: pass others
  | [Tree left; Elt root; Tree right] → [join left root right]
  | _ → assert false

let rec loop = function
  | [] → Leaf
  | [Tree t] → t
  | list → loop (pass list)
```

Notice how the invariant that alternating lists have length $2^k - 1$ is maintained: this is because, for $k \geqslant 2$, we have $2^k - 1 = 4 \times (2^{k-2} - 1) + 3$, hence we obtain an alternating list of length $2 \times (2^{k-2} - 1) + 1 = 2^{k-1} - 1$.

It remains to show how to transform a list of labels of length n into an alternating list of trees and labels. Each of the original trees has height zero or one: they are leaves or contain only one label. Because we want a list of length precisely $2^k - 1$, for $k = 1 + \lfloor \log_2 n \rfloor$, it means we need $2^k - 1 - n$ leaves. This quantity is computed as the variable missing. The function pad computes the alternating list by creating as many leaves as needed, alternating them with elements, and once enough leaves are created, promotes all the odd-position labels into trees.

```
let complete list =
  let n = List.length list in
  let rec pow2 i = if i <= n then pow2 (2*i) else i in
  let missing = (pow2 1) - n - 1 in
  let rec pad missing = function
    | head::tail when missing <> 0 →
      Tree Leaf :: Elt head :: pad (missing - 1) tail
    | odd::even::others → join Leaf odd Leaf :: Elt even :: pad 0 others
    | [single] → [join Leaf single Leaf]
    | [] → []
  in
  pad missing list
```

The balancing algorithm balance: α list $\to$ α tree is thus given by the composition of loop with complete:

```
let balance list = loop (complete list)
```

As for the complexity of this algorithm, notice that pass and complete are both clearly in linear-time in the length of the lists on which they work, while loop recurses on lists whose length are halved at each recursive step. Hence balance is a linear-time algorithm.

3 Removing partial functions

The loop function of Section 2 relies on the invariant that the list argument has length $2^k - 1$. Additionally, all the odd-position values must be of the form Tree t, whereas all the even-position values must be of the form Elt x. If either of these invariants is broken, we would run into the assert false of pass.

It is not immediately apparent that these properties hold. If it does not take a tremendous effort to convince oneself that the balance function of Section 2 is indeed correct, a direct mechanically checked proof would not be very practical.

3.1 Length invariants

Our goal in this section is to avoid resorting to assert false. In addition to making sure that balance indeed terminates with a value, it will make it considerably simpler to implement the balancing algorithm in Coq in Section 4. To achieve this goal, it is necessary to have more precise types. Let us focus first on the length invariants: we will need to define a type which contains exactly the non-empty lists of length $2^k - 1$.

A data structure which holds $2^k - 1$ elements brings complete binary trees to mind. Even if it is possible – though not necessary convenient – to represent complete binary trees in Ocaml, they are not the appropriate structure. First, because complete binary trees are full trees and are, hence, unlikely to serve as a useful intermediate data structure to build a full tree. Second because there is a simpler – albeit more exotic – alternative.

Indeed, lists can be seen as decorated unary numbers: there is an element at each successor. Different kinds of lists can be obtained, more or less systematically, by varying the numerical representation. This idea goes back to Guibas *&* al. in [1] and a fairly thorough exploration can be found in Okasaki [2, Chapters 9*&*10]. In the simplest cases, the analogous list structure corresponds to a structurally recursive exponentiation algorithm. For regular lists, a list of size n whose elements have type a can be recursively defined with the following equations:

$$\begin{cases} a^0 & = 1 \\ a^{n+1} & = a \times a^n \end{cases}$$

Replacing unary numbers with binary numbers, we obtain the binary exponentiation algorithm:

$$\begin{cases} a^{2^0-1} & = 1 \\ a^{2n} & = (a^2)^n \\ a^{2n+1} & = a \times (a^2)^n \end{cases}$$

Okasaki [2, Chapter 10] uses a non-uniform inductive type to encode the latter exponentiation algorithm into a type of lists he calls *binary lists*. We are only interested in lists of length 2^k-1, that is a length written only with the digit 1 in binary representation. So following Okasaki, but skipping the second equation above (which corresponds to the digit 0) we define the following non-uniform inductive type, which we call *power lists*:

```
module PowerList = struct

  type α t =
    | Zero
    | TwicePlusOne of α * (α*α) t

end
```

This type actually appears in Okasaki [2, Chapter 10] as an introduction to non-uniform binary lists. Relatedly, Okasaki [3] leverages a tail-recursive binary exponentiation algorithm to define a type capturing precisely square matrices;

on the other hand, Myers [4] introduced a flavour of list based on *skew binary numbers* which are not easily captured as exponentiation.

Although the power lists may look like some sort of trees, it is not a very accurate depiction. The easiest way to picture how power lists works is to see TwicePlusOne as a fancy `(::)`, then the lists with, respectively, 1, 3, 7, and 15 elements are as follows:

- `[1]`
- `[1;(2,3)]`
- `[1;(2,3);((4,5),(6,7))]`
- `[1;(2,3);((4,5),(6,7));(((8,9),(10,11)),((12,13),(14,15)))]`

Elements appear in order, like in a regular list, but they are packed twice as tightly after each TwicePlusOne.

Just like with regular lists, there is a *map* function for power lists. Due to the non-uniformity it is a little trickier[3] than the regular list map: in the recursive steps, the argument function f needs to process two consecutive elements at a time.

```
module PowerList = struct
  ⋮
 let rec map : α β. (α→β) → α t → β t = fun f → function
   | Zero → Zero
   | TwicePlusOne (elt,lst) →
     let f' (x,y) = f x , f y in
     TwicePlusOne (f elt , map f' lst)

end
```

3.2 Alternation

In Section 2, labels are separated from trees dynamically. The pass function verifies that trees and labels are interleaved properly, and fails if they are not.

In this section, instead, we consider a variant of α PowerList.t where every odd position contains a tree, and every even position contains an element. More generally, we define a type `(ω,η) AlternatingPowerList.t` where odd positions have type ω, and even positions have type η. Such a list should have the following pattern:

- `[ω;(η,ω);((η,ω),(η,ω))]`

After the first element, which must have type ω, there is no difference between even and odd positions: indeed, excluding the first element, we are actually

[3] The type annotation on PowerList.map informs Ocaml that map is a non-uniform recursive function. Without the type annotation, Ocaml simply assumes that map is uniformly recursive and fails to typecheck since f and f' have different types.

building an (η*ω) PowerList.t. Hence the definition:

```
module AlternatingPowerList = struct

  type (ω,η) t =
    | Zero
    | TwicePlusOne of ω * (η*ω) PowerList.t

end
```

For brevity, let us write PL and APL for PowerList and AlternatingPowerList respectively.

Using these alternating power lists, we can define a version of the pass function free of assert false. Indeed, consider an alternating power list of length at least 3: it is of the form APL.TwicePlusOne (a, PL.TwicePlusOne ((b,c), lst)), where lst has type ((η*ω)*(η*ω)) PowerList.t. The pass function of Section 2, as it happens, manipulates its arguments by groups of four elements: basically, pass is simply a map over lst.

We hence define the function pass which joins the trees in a list of length $2^{k+2} - 1$[4], producing a list of length $2^{k+1} - 1$. The function loop is virtually unchanged from Section 2, except it acts on power lists:

```
let pass left (root,right) apl =
  let join_up ((single,left),(root,right)) =
    single, Node (left,root,right)
  in
  APL.TwicePlusOne ( Node (left,root,right) , PL.map join_up apl)

let rec loop : ϵ. (ϵ tree,ϵ) APL.t → ϵ tree = function
  | APL.Zero → Leaf
  | APL.TwicePlusOne (tree,PL.Zero) → tree
  | APL.TwicePlusOne (tree,PL.TwicePlusOne (pair,apl)) →
      loop (pass tree pair apl)
```

3.3 Padding

Now that there is no more assert false in the code of loop, we need to change the complete function of Section 2 so that it returns an (α tree,α) APL.t rather than a list. The heart of this section is a function which turns an α list into an (α*α tree) PL.t. The final function, which produces an (α tree,α) APL.t is a simple wrapper around the former.

We want to turn a list lst of length $n + 1$ into a pair of its first element, converted into a tree, plus a power list of length $2 \times (2^k - 1) \geqslant n$ representing its tail tail. Each element of the power list is a pair, whose first term is an element, and its second term is a tree of height at most one. In particular, the length of

[4] To ensure that its argument list has at least three elements, pass takes the three first elements as extra arguments. In other words pass t (x,s) l is meant to be read as pass (APL.TwicePlusOne (t , PL.TwicePlusOne ((x,s),l))).

the returned power list is always even, so if tail has odd length, we will need to insert at least a Leaf. This suggests that we may inspect the parity of the length of tail, and insert an extra element precisely when it is odd. This leads to a slightly different padding procedure than that of Section 2, in particular the leaves are not inserted at the same position, but it is inconsequential.

An α list of even length can be turned into an (α*α) list whose length is halved. This turns out to be interesting for our recursion, since it mimics the inductive step of power lists. Also, in the case of even length, we need to distinguish the empty case from the non-empty case: the former will be turned into the empty power list APL.Zero while the latter will be turned into a power list of the form APL.TwicePlusOne((x,y),l), where x and y correspond to the two first elements of tail. These different cases are represented in the following view:

```
type α parity =
  | Empty
  | Odd of α * (α*α) list
  | Even of (α*α) * (α*α) list

let pair_up lst =
  let succ elt = function
    | Empty → Odd (elt, [])
    | Odd (b,pairs) → Even ((elt,b), pairs)
    | Even (bc,pairs) → Odd (elt, bc::pairs)
  in
  List.fold_right succ lst Empty
```

The padding function itself, of_list, is at first sight far from intuitive. Let us recall that we want to turn a list of labels of arbitrary length, into a power list of pairs. A label can be thought of as a bit of weight 2^0, while a pair of labels would be a bit of weight 2^1, and so on. At first, all our bits have weight 2^0 and consists in one label each. We can build bits of higher weight by pairing up two bits of the same weight. A bit made up only of labels is called *pure*. We can also double the weight of a bit by interlacing leaves with it (with the function pad), but this gives a bit made of pairs of labels and trees, call them *impure*. Lastly, we can also transmute a pure bit into an impure bit of the same weight (with the function coerce), by replacing odd-position labels by trees of height one.

Each recursive step consists in taking a list of pure bits of the same weight 2^k, and outputing exactly one impure bit of size 2^{k+1}, plus a list of pure bits of weight 2^{k+1}, which is converted recursively. We thus obtain, successively, one bit of each weight from 2^1 to 2^l, for some l, encoding a list of length $2^{l+1} - 2$, as expected.

At any recursive step, suppose first that the number of bits of weight 2^k is odd. As we need to compute only bits of weight 2^{k+1}, one of them impure, we are forced to use pad on one bit, and to pair up the others. Suppose now that the number of bits of weight 2^k is even. In that case, we can pair them all into bits of weight 2^{k+1}, and then use coerce on one of them to make the impure bit.

The last difficulty is that pad and coerce both depend on the current weight of the bits, hence we need to update them at each recursive step. pad must add leaves between every two consecutive labels, in even positions, while coerce must upgrade every even-position label into a tree of height one. This leads to the following definition:

```
module PowerList = struct
  ⋮
 let rec of_list : α β. (α→β) → (α*α→β) → α list → β t =
  fun pad coerce bits →
   let pad' (x,y) = (pad x, pad y) in
   let coerce' (x,y) = (coerce x, coerce y) in
   match pair_up bits with
   | Empty → Zero
   | Odd (a,pures) → TwicePlusOne (pad a, of_list pad' coerce' pures)
   | Even (ab,pures) →
     TwicePlusOne (coerce ab, of_list pad' coerce' pures)

end
```

With that function, we can conclude our implementation. Again writing PL and APL for PowerList and AlternatingPowerList respectively:

```
module AlternatingPowerList = struct
  ⋮
 let of_list leaf up id = function
   | [] → Zero
   | a::l →
     let pad x = id x , leaf in
     let coerce (x,y) = id x , up y in
     TwicePlusOne (up a, PowerList.of_list pad coerce l)

end

let singleton x = Node(Leaf,x,Leaf)
let balance l =
 loop (APL.of_list Leaf singleton (fun e→e) l)
```

The final function, balance : α list → α tree, implements the same algorithm as Section 2 without any partial functions.

What we may have lost in this section, compared to the simple algorithm, is the simplicity of the complexity analysis of the algorithm. The subtleties of the main functions require a finer analysis. Consider first the function PL.map: clearly the number of recursive calls depends only logarithmically on the number of elements in the power list. But each recursive call uses as its first argument a function twice as complex than the previous one. This leads to the following inequation over the complexity $C(n, m, f)$ of map, where n is the number of

elements in the power list, m is the size of the elements in the power list, and f is the complexity of the mapped function:

$$C(n, m, f) \leq f(m) + C\left(\frac{n-1}{2}, 2m, k \mapsto 2 \times f(k/2) + O(1)\right) + O(1)$$

From there, it is easy to prove that $C(n, m, f) = n.f(m) + O(n)$, so that `PL.map` runs indeed in linear-time, and so is `loop`. Similarly the complexity of `PL.of_list` can be described by a higher-order recursive inequation (almost the same as above, except that the complexity depends on two functions and the constant term is replaced by a linear term), whose solution gives also a linear-time complexity.

4 Turning to Coq

There is still a property of the algorithm that the implementation of Section 3 does not make obvious: that the algorithm actually does build *full* trees. In this section we shall build into the type of `balance` that its output is indeed full.

To that effect, we will use Coq rather than Ocaml. Even if it is possible, with some effort, to represent full trees and implement the algorithm in Ocaml – and relatively easy in Haskell – a Coq implementation also gives us termination by construction. Coq forces every recursion to be structural, which will prove to be rather entertaining.

At a superficial level, a visible difference with the Ocaml implementation is that `Powerlist.t` and `AlternatingPowerList.t` must be decorated with the k such that the length is $2^k - 1$: it is the structural recursion parameter of the `balance_powerlist` function. Because it makes the code simpler, we will use a recursive definition rather than an inductive one:

```
Module PowerList.

 Fixpoint T (A:Type) (k:nat) :=
  match k with
  | 0 ⇒ unit:Type
  | S k' ⇒ A * T (A*A) k'
  end.

End PowerList.
```

We will also need a version where `k` can be arbitrary. For that purpose we use Coq's type of dependent pairs `{ n:nat & F n}`. The constructor for dependent pairs is written `⟨ n , x ⟩`. The implicit version comes with constructors – `tpo`

stands for "twice plus one":

```
Module PowerList.
  ⋮
 Definition U (A:Type) ≔ { k:nat & T A k }.
 Definition zero {A:Type} : U A ≔ ⟨ 0 , tt ⟩.
 Definition tpo {A:Type} (a:A) (l:U (A∗A)) : U A ≔
  let '⟨k,l⟩ ≔ l in
  ⟨ S k , (a,l) ⟩.
End PowerList
```

The definition of AlternatingPowerList.T and AlternatingPowerList.U are similar.

4.1 Full trees

To code full trees, we index trees by their height, and specify that leaves can happen only at height 0 or 1:

```
Inductive FullTree (A:Type) : nat → Type ≔
 | Leaf₀ : FullTree A 0
 | Leaf₁ : FullTree A 1
 | Node {k:nat} : FullTree A k → A → FullTree A k → FullTree A (S k).
```

If we omitted the constructor Leaf_1, we would have a definition of complete binary trees: both subtrees of a node are complete binary trees of the same height. We allow the full trees to be incomplete by letting Leaf_1 take the place of nodes on the last level.

Using the type FullTree A k in place of the type α tree, the functions pass and balance_powerlist are virtually unmodified[5] with respect to Section 3. Only their types change to reflect the extra information:

```
Definition pass {A k p} : APL.T (FullTree A (S p)) A (S (S k)) →
    APL.T (FullTree A (S (S p))) A (S k).
Fixpoint loop {A k p} : APL.T (FullTree A (S p)) A (S k) →
    FullTree A (plus k (S p)) {struct k}.
```

The algorithm indeed builds only full trees.

4.2 Structural initialisation

The padding conversion from lists to power lists, in Section 3, is not structural due to the use of pair_up in the recursive call. To tackle this recursion, we shall

[5] In fact, as can be seen from its type, loop only handles non-empty alternating power lists. This is due to a small technicality: the recursive step of loop is the case S (S k), but Coq does not recognise S k as a structural subterm of S (S k), so the definition from Section 3 does not verifies Coq's structural recursion criterion. As a workaround, the empty case is moved to the balance function.

make use of another intermediate structure. What we need, essentially, is that all the calls to pair_up are pre-calculated, so the intermediate structure will be like parity except that the calls to (α∗α) list are replaced by inductive calls.

As it turns out, this is another non-uniform datatype which corresponds to a numerical representation. Indeed, any natural number can be written in binary with digits 1 and 2 (but not 0). In this system, for example, $8 = 1 \times 2^2 + 1 \times 2^1 + 2 \times 2^0$ is represented as 112. Here is the definition, where tpo reads "twice plus one" and tpt "twice plus two":

```
Module BinaryList.

 Inductive T (A:Type) : Type :=
 | zero
 | tpo (a:A) (l:T (A∗A))
 | tpt (a b: A) (l:T (A∗A)).

End BinaryList.
```

To turn a non-empty list into a BinaryList.T, all we need is a function cons of type A → T A → T A to add an element in front of the list. On the numerical representation side, it corresponds to adding 1. It behaves like adding 1 in the usual binary representation, except that 1-s are turned into 2-s without a carry and 2-s into 1-s while producing a carry:

```
Module BinaryList.
   ⋮
 Fixpoint cons {A} (a:A) (l:T A) : T A :=
  match l with
  | zero ⇒ tpo a zero
  | tpo b l ⇒ tpt a b l
  | tpt b c l ⇒ tpo a (cons (b,c) l)
  end.

 Definition of_list {A} (l:list A) : T A :=
  List.fold_right cons zero l.

End BinaryList.
```

Note that while cons takes, in the worst case, logarithmic time with respect to the length of the list, building a list by repeatedly using cons is still linear. Indeed, as previously mentioned, cons mimics the successor algorithm for binary numbers, whose amortized complexity is well-known to be constant.

We also need a function which turns a T (A∗A) into a T A. This is effectively multiplication by 2. The lack of 0 among the digits[6] makes this process recursive. A simple presentation of the doubling algorithm consists in adding a 0 at the

[6] The constructor zero represents an empty list of digits.

end of the number, then eliminating the 0 using the following equalities:

$$\begin{cases} 0 = \cdot \\ x20 = x12 \\ x10 = x02 \end{cases}$$

In terms of binary lists:

```
Module BinaryList.
  ⋮
 Fixpoint twice {A} (l:T (A*A)) : T A :=
  match l with
  | zero ⇒ zero
  | tpo (a,b) l ⇒ tpt a b (twice l)
  | tpt (a,b) cd l ⇒ tpt a b (tpo cd l)
  end.

End BinaryList.
```

We can now write a structurally recursive padding function, using binary lists as the structural argument. As we do not know in advance the length of the produced list, a PowerList.U is returned. We write BL as a shorthand for BinaryList:

```
Module PowerList.
  ⋮
 Fixpoint of_binary_list {A X} (d:A→X) (f:A*A→X) (l:BL.T A) : U X :=
  match l with
  | BL.zero ⇒ zero
  | BL.tpo a l ⇒
   tpo (d a) (of_binary_list (d×d) (f×f) l)
  | BL.tpt a b l ⇒
   tpo (f (a,b)) (of_binary_list (d×d) (f×f) l)
  end.

End PowerList.
```

Where g×f is the function which maps (x,y) to (g x,f y).

The rest follows straightforwardly, and we can define the following functions which conclude the algorithm (BL, PL, and APL stand for BinaryList, PowerList,

and AlternatingPowerList respectively):

```
Module AlternatingPowerList.
  ⋮
 Definition of_binary_list {A Odd Even}
    (d:Odd) (f:A→Odd) (g:A→Even) (l:BL.T A) : U Odd Even :=
  match l with
  | BL.zero ⇒ zero
  | BL.tpo a l ⇒
    let d' x := ( g x , d ) in
    tpo (f a) (PL.of_binary_list d' (g×f) (BL.twice l))
  | BL.tpt a b l ⇒
    let d' x := ( g x , d ) in
    tpo (f a) (PL.of_binary_list d' (g×f) (BL.tpo b l))
  end.

 Definition of_list {A Odd Even}
    (d:Odd) (f:A→Odd) (g:A→Even) (l:list A) : U Odd Even :=
  of_binary_list d f g (BL.of_list l).

End AlternatingPowerList.

Definition singleton {A:Type} (x:A) : FullTree A 1 :=
 Node Leaf_0 x Leaf_0.

Definition balance {A:Type} (l:list A) : { k:nat & FullTree A k } :=
 let '⟨k,l⟩ := APL.of_list Leaf_1 singleton (fun x⇒x) l in
 match k with
 | 0 ⇒ fun _ ⇒ ⟨ 0 , Leaf_0 ⟩
 | S k ⇒ fun l ⇒ ⟨ plus k 1 , loop l ⟩
 end l.
```

5 Conclusion

The balance function of Section 4 is, by virtue of its type alone, a total function which turns lists into full binary trees. Yet, to the cost of using intermediary data-structures, it effectively implements the algorithm of Section 2.

The missing piece is to prove that the infix traversal of balance l is indeed l. The infix traversal of a (full) tree is represented in Coq with the functions

```
Fixpoint list_of_full_tree_n {A n} (t:FullTree A n) : list A :=
  match t with
  | Leaf_0 ⇒ []
  | Leaf_1 ⇒ []
  | Node _ t_1 x t_2 ⇒
    list_of_full_tree_n t_1 ++ [x] ++ list_of_full_tree_n t_2
  end.

Definition list_of_full_tree {A} (t:{ k:nat & FullTree A k }) : list A :=
  list_of_full_tree_n (projT2 t).
```

We can then state the theorem:

```
Theorem balance_preserves_order A (l:list A) :
  list_of_full_tree (balance l) = l.
```

The proof is short and straightforward: we define a traversal function for each intermediate structure; and state a variant of balance_preserves_order for each intermediate function. Proving the intermediate lemmas is not difficult and can be mostly automatised: we use a very simple generic automated tactic, which discharges most goals. This theorem concludes our easy formal proof of the balancing algorithm.

References

1. Guibas, L.J., McCreight, E.M., Plass, M.F., Roberts, J.R.: A new representation for linear lists. In: Proceedings of the ninth annual ACM symposium on Theory of computing - STOC '77. STOC '77, New York, New York, USA, ACM Press (1977) 49–60
2. Okasaki, C.: Purely functional data structures. Cambridge University Press (1999)
3. Okasaki, C.: From fast exponentiation to square matrices. In: Proceedings of the fourth ACM SIGPLAN international conference on Functional programming - ICFP '99, New York, New York, USA, ACM Press (1999) 28–35
4. Myers, E.W.: An applicative random-access stack. Information processing letters **17** (1983) 241–248